

# EXHIBIT O

UNITED STATES DISTRICT COURT  
FOR THE EASTERN DISTRICT OF TEXAS  
MARSHALL DIVISION

NETLIST, INC, )  
                  )  
Plaintiff,     )  
vs.              )  
SAMSUNG ELECTRONICS CO., LTD.,     ) Civil Action No. 2:21-CV-463-JRG  
SAMSUNG ELECTRONICS AMERICA,     )  
INC., SAMSUNG SEMICONDUCTOR,     )  
INC.,             )  
                  )  
Defendants.     ) JURY TRIAL DEMANDED

---

**PLAINTIFF NETLIST, INC.'S OPENING CLAIM CONSTRUCTION BRIEF**

## TABLE OF CONTENTS

	<u>Page</u>
I. Introduction .....	1
II. The '339 Patent.....	1
A. Samsung's Attempt to Inject the "Fork-In-The-Road" Configuration into the Plain Language of the Claims Should Be Rejected ('339 claims 1, 11, 19, 27) .....	2
1. Background: "Fork-In-the-Road" vs. "Straight Line" Configurations .....	3
2. The Claims of the '339 Patent Are Not Limited to "Fork-in- the-Road" .....	5
B. The Claims Do Not Require Rank Multiplication ('339 claims 1, 11, 19, 27) .....	8
C. The Court Should Resolve the Parties' Dispute over the Plain Meaning of "time period in accordance with a latency parameter" .....	12
III. The '506 Patent.....	15
A. Claim 14 of the '506 Patent Is Not Limited to "determining a first predetermined amount" During "one or more previous memory operations" .....	16
IV. The PMIC Patents ('918/'054).....	17
A. The Multiple "regulated voltages" And "voltage amplitude" Need Not Have Different Voltage Levels.....	18
B. Samsung's Proposed Construction For "A Second Plurality of Address And Control Signals" Confuses Rather Than Clarifies the Meaning of the Term .....	22
C. A "dual buck converter" Need Not Produce Two "Distinct" Voltages.....	23
D. The Preamble of the PMIC Patent Claims Is Limiting.....	24
E. "pre-regulated input voltage" / "input voltage" .....	25
V. The HBM Patents ('060/'160).....	26
A. "array die" Does Not Exclude a DRAM Circuit .....	27

	<u>Page</u>
B. A “chip select signal” Is Not Limited To Selecting Only One Array Die.....	30

## **TABLE OF AUTHORITIES**

	Page(s)
<b>Cases</b>	
<i>Apple, Inc. v. Ameranth, Inc.</i> , 842 F.3d 1229 (Fed. Cir. 2016).....	20
<i>Catalina Mktg. Int'l, Inc. v. Coolsavings.com, Inc.</i> , 289 F.3d 801 (Fed. Cir. 2002).....	24
<i>Genuine Enabling Tech. LLC v. Nintendo Co.</i> , 29 F.4th 1365 (Fed. Cir. 2022) .....	28
<i>Hill-Rom Services, Inc. v. Stryker Corp.</i> , 755 F.3d 1367 (Fed. Cir. 2014).....	6, 21, 24
<i>Kaufman v. Microsoft Corp.</i> , 34 F.4th 1360 (Fed. Cir. 2022) .....	11
<i>Oatey Co. v. IPS Corp.</i> , 514 F.3d 1271 (Fed. Cir. 2008).....	11
<i>Pbillips v. AWH Corp.</i> , 415 F.3d 1303 (Fed. Cir. 2005).....	20
<i>Tech. Props. Ltd. LLC v. Huawei Techs. Co.</i> , 849 F.3d 1349 (Fed. Cir. 2017).....	28
<i>In re Varma</i> , 816 F.3d 1352 (Fed. Cir. 2016).....	28
<i>Virnetx, Inc. v. Cisco Sys., Inc.</i> , 767 F.3d 1308 (Fed. Cir. 2014).....	17
<b>Other Authorities</b>	
IEEE 100 The Authoritative Dictionary of IEEE Standards Terms (7th Ed. 2000).....	27
Webster's New World, College Dictionary (4th Ed., 2000) .....	15, 19
Wiley Electrical and Electronics Engineering Dictionary (2004) .....	15, 19

## I. Introduction<sup>1</sup>

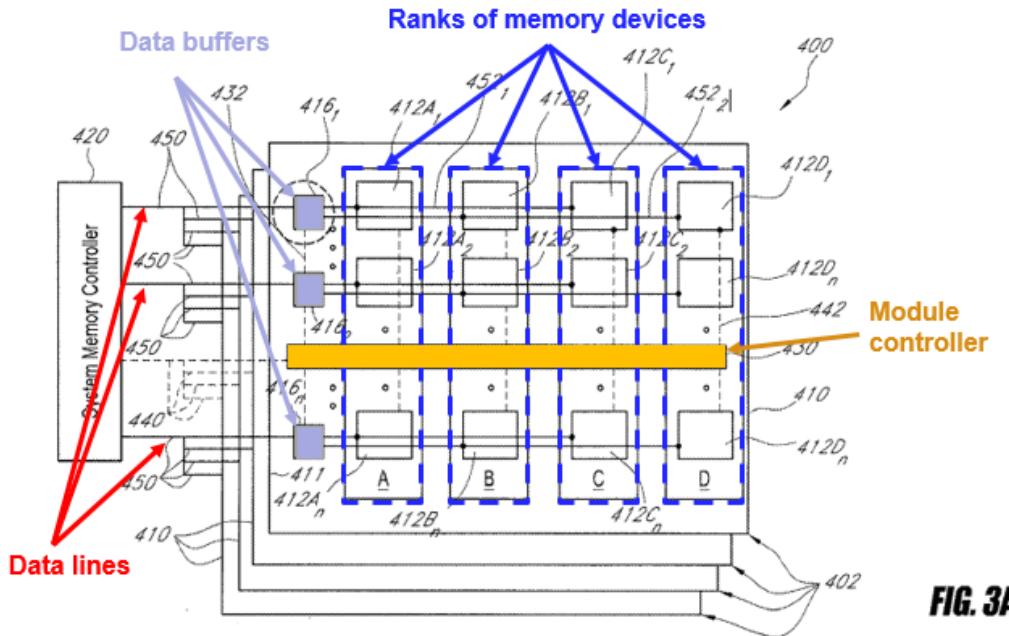
Plaintiff, Netlist, Inc. (“Netlist”) respectfully submits this opening claim construction brief on U.S. Patent Nos. 10,949,339 (the “‘339 patent”), 10,860,506 (the “‘506 patent”), 11,016,918 (the “‘918 patent”), 11,232,054 (the “‘054 patent” and together with the ‘918 patent, the “PMIC patents”), 8,787,060 (the “‘060 patent”), and 9,318,160 (the “‘160 patent” and together with the ‘060 patent, the “HBM patents”). *See* Exs. 1-6. Netlist’s positions are supported by the patents’ intrinsic evidence and further confirmed by extrinsic evidence.

## II. The ‘339 Patent

The ‘339 patent is directed to specific features related to the control and timing of data signals in a memory module featuring a distributed buffer architecture, like the accused load-reduced dual-in-line memory modules (“LRDIMMs”). The ‘339 patent generally discloses a memory module configured to communicate with a memory controller of the system (420). The memory module includes double data rate (“DDR”) DRAM devices arranged in multiple ranks, each rank of the same width as the memory module. The module also includes a module controller (430) configured to receive input address and control signals for a read or write operation from the host’s memory controller and to output registered address and control signals. Ex. 1 (‘339), 1:24-43, 2:26-61. A host can access the memory modules via the memory controller in the host computer system, which transmits commands like read or write. *Id.*, 1:24-36, 1:52-59. As the number of memory devices on a module increases, so too does the load presented by the data lines, resulting in higher power dissipation, slower operational speed, and higher costs. *Id.*, 2:5-12, 4:27-47. To reduce the load presented to the memory controller, data buffers 416 (aka “data transmission circuits”) are placed between the data lines 450 from the memory controller and the memory devices 412. *Id.*, 10:66-11:4.

---

<sup>1</sup> All emphasis is added unless stated otherwise.



**FIG. 3A**

Accommodating data buffers in conventional memory modules required precise timing of the data going through the data path in the data buffer. *Id.*, 5:23-43. The '339 patent solves this problem by, among other things, using a module controller that coordinates the transmission of control signals to both the memory devices and the data buffers. *Id.*, 10:17-32, 10:33-53. In response to the module control signals, logic in the data buffer enables tristate buffers and the data path in the data buffer for a time period in accordance with a latency parameter to actively drive a byte-wise section of the data from one side of the data buffer to the other side during the time period. *See, e.g., id.*, 2:47-61, 3:26-28, 14:46-58, 15:61-16:44, 17:45-18:65, 19:53-67, 22:16-22:30.

The disputed terms of the '339 patent relate to the arrangement of DRAMs in relation to the data buffers, whether the module controller must be configured to support a feature called rank multiplication, and the data-transmission timing issue.

**A. Samsung's Attempt to Inject the "Fork-In-The-Road" Configuration into the Plain Language of the Claims Should Be Rejected ('339 claims 1, 11, 19, 27)**

Claim 1, and Samsung's proposed construction, is representative of the parties' dispute.

Term	Netlist	Samsung
“each respective byte-wise buffer further includes logic configurable to <u>control the byte-wise data path</u> in response to the module control signals, wherein <u>the byte-wise data path is enabled for a first time period</u> in accordance with a latency parameter to actively drive a respective byte-wise section of the N-bit wide write data associated with the memory operation from the first side to the second side during the first time period” <sup>2</sup> (339, claim 1)	“time period in accordance with a latency parameter” means “a time period wherein both the start of the time period and duration of the time period depends on at least a latency parameter” (see Section C)	“each respective byte-wise buffer further includes logic configurable to, in response to the module control signals, activate the byte-wise data path connected to a first DDR DRAM device (in a first N-bit-wide rank), <i>and disable the byte-wise data path connected to a second DDR DRAM device (in a second N-bit-wide rank)</i> , to cause a respective byte-wise section of the N-bit wide write data associated with the memory operation to be sent from the first side to the first DDR DRAM device along the activated byte-wise data path and <i>not sent to the second DDR DRAM device along the disabled byte-wise data path</i> during the first time period in accordance with a latency parameter.”
“logic in response to the module control signals is configured to <u>enable the first tristate buffers to drive the respective byte-wise section of the N-bit wide write data</u> to the respective module data lines during the first time period” (339, claim 1)	No construction is necessary (i.e., plain and ordinary meaning).	“logic in response to the module control signals is configured to enable the first tristate buffers to activate the byte-wise data path connected to a first DDR DRAM device (in a first N-bit-wide rank) <i>and disable the byte-wise data path connected to a second DDR DRAM device (in a second N-bit-wide rank)</i> , to cause the respective byte-wise section of the N-bit wide write data to be sent from the first side to the module data lines coupled to the first DDR DRAM device along the activated byte-wise data path and <i>not sent to the module data lines coupled to the second DDR DRAM device along the disabled byte-wise data path</i> during the first time period”

### 1. Background: “Fork-In-the-Road” vs. “Straight Line” Configurations

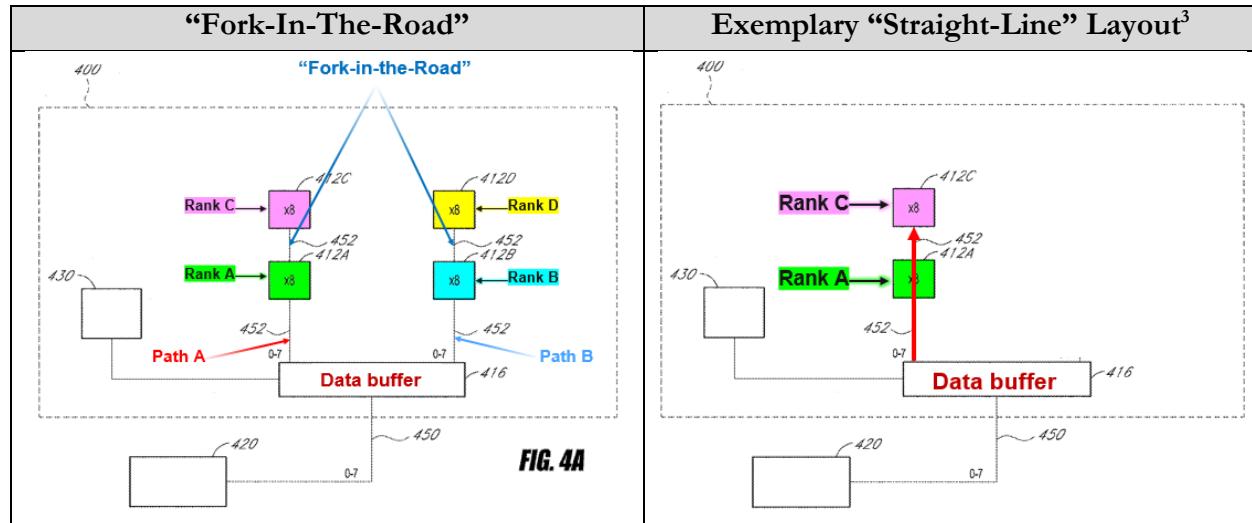
In previous litigation against SK hynix involving U.S. Patent Nos. 8,516,185 and 9,606,907, the parties used the terms “Straight-Line” and “Fork-in-the-Road” to refer to configurations for controlling transmission of data to and from the memory devices on the memory module through the

---

<sup>2</sup> This clause combines two of Samsung’s identified terms, per its suggestion. Dkt. 70-2 at 32 n.3.

data transmission circuits.

In the “Fork-in-the-Road” arrangement, not all ranks of memory devices connected to a data buffer share a common data path. Ex. 1 ('339), 11:4-11:44, Fig. 3A (ranks A and C share one path while ranks B and D share a second path, i.e. **separate** data paths); Ex. 10 (SS '339 IPR Pet.) at 13 (annotating Fig. 5 to illustrate “Fork-in-the-Road” arrangement); *see also* illustration below.



In a “Fork-in-the-Road” configuration, switching between the multiple data paths may occur. Ex. 1 ('339), 11:35-44 (“In certain embodiments, at least one data transmission circuit 416, 416' selectively switches between two or more memory devices 412, 412' so as to operatively couple at least one selected memory device 412, 412' to the system memory controller 420, 420.”); *see also id.*, Fig. 4A (annotated, left). In the “Straight Line” arrangement, ranks of memory devices are on the **same data path** without any “Fork in the Road.” *See* modified Fig. 4A (annotated, right). In this arrangement, the buffer allows a single load to be presented to the memory controller regardless of the number of memory devices that are on the path. *Id.*, 15:13-18, 16:45-54.

In prior administrative proceedings, the ITC determined whether a “Fork-in-the-Road” configuration should apply based on the claim language. For instance, in an ITC investigation

---

<sup>3</sup> Modified Figure 4A with two ranks of memory devices connected to the data buffer.

involving the related '185 patent, ALJ Bullock found that a “Fork-in-the-Road” layout applied only because the claims required that the memory devices be “selectively isolated” by the claimed buffer circuits, and that these buffer circuits “selectively allow” data transactions. Ex. 7 (11/14/2017 1023-ID) at 120-121 (“‘selectively isolate’ conveys that during a write operation a ‘selected’ subset of the memory modules are electrically separated from the memory controller thereby reducing the load upon the memory controller even though the ‘electrically isolated’ memory modules are still ‘operatively coupled’ to the memory controller.”). In contrast, in another ITC proceeding involving the parent '907 patent, ALJ Pender ultimately agreed with Netlist that the “Fork-in-the-Road” configuration did not apply to the claims of the '907 patent. *See* Ex. 8 (8/30/2018 Claim Construction Order) at 27-31; Ex. 9 (4/21/2020 1089-Opinion) at 8-9.<sup>4</sup>

## 2. The Claims of the '339 Patent Are Not Limited to “Fork-in-the-Road”

Samsung’s attempt to inject the “Fork-in-the-Road” restriction into the claims of the '339 patent is improper. First, the claims of the '339 patent do not include the language that ALJ Bullock interpreted as limited to “Fork-in-the-Road” in the 2017 ITC litigation. For example, unlike the '185 patent, the claims of the '339 patent do not require that the buffer circuits be configured to “selectively isolate” the non-targeted memory devices or to “selectively allow” data transmissions:

'339 patent, claim 1	'185 patent (Ex. 11), claim 1
“each respective byte-wise buffer further includes logic configurable to control the byte-wise data path in response to the module control signals, wherein the byte-wise data path is enabled for a first time period in accordance with a latency parameter to actively drive a respective byte-wise section of the N-bit wide write data associated with the memory operation from the first side to the second side during the first time period”	“each circuit of the plurality of circuits ... configured <b>to selectively allow</b> data transmission between the system memory controller and at least one selected memory device of the at least two corresponding memory devices in response to the module control signals, and to <b>selectively isolate at least one other memory device of the at least two corresponding memory devices from the system memory controller in response to the module control signals</b> ”

<sup>4</sup> The Commission later decided to resolve the parties’ dispute based on the plain meaning of the term “receive” in the phrase “do not output or receive” instead. Ex. 9 (4/21/2020 1089-Opinion) at 10. The claims at issue here do not have the same “receive” language.

Instead, like the claims of '907 patent that were found not subject to the "Fork-in-the-Road" configuration, the claims at issue also do not have terms that require the "Fork-in-the-Road" restriction.<sup>5</sup> Moreover, the claims of the '339 patent do not contain any requirement that the "second memory devices" (which may be a device in a "second N-bit wide rank") be configured to "not output or receive data" during a particular transaction, which was what precipitated the dispute over whether the "Fork-in-the-Road" configuration applied in the context of the '907 patent. Ex. 12 ('907), claim 1; Ex. 8 at 24-25 (noting that the term "output or receive data / do not output or receive any data" was "at the heart of the larger dispute between the parties as to whether the asserted claims of the '907 patent require a 'fork in the road' or 'straight line' arrangement of memory devices").

Second, the inventors never uttered any "words or expressions of manifest exclusion or restriction" to limit the plain claim language to a "Fork-in-the-Road" configuration. *Hill-Rom Services, Inc. v. Stryker Corp.*, 755 F.3d 1367, 1371-72 (Fed. Cir. 2014) (citations omitted). To the contrary, the specification expressly contemplates a straight-line configuration:

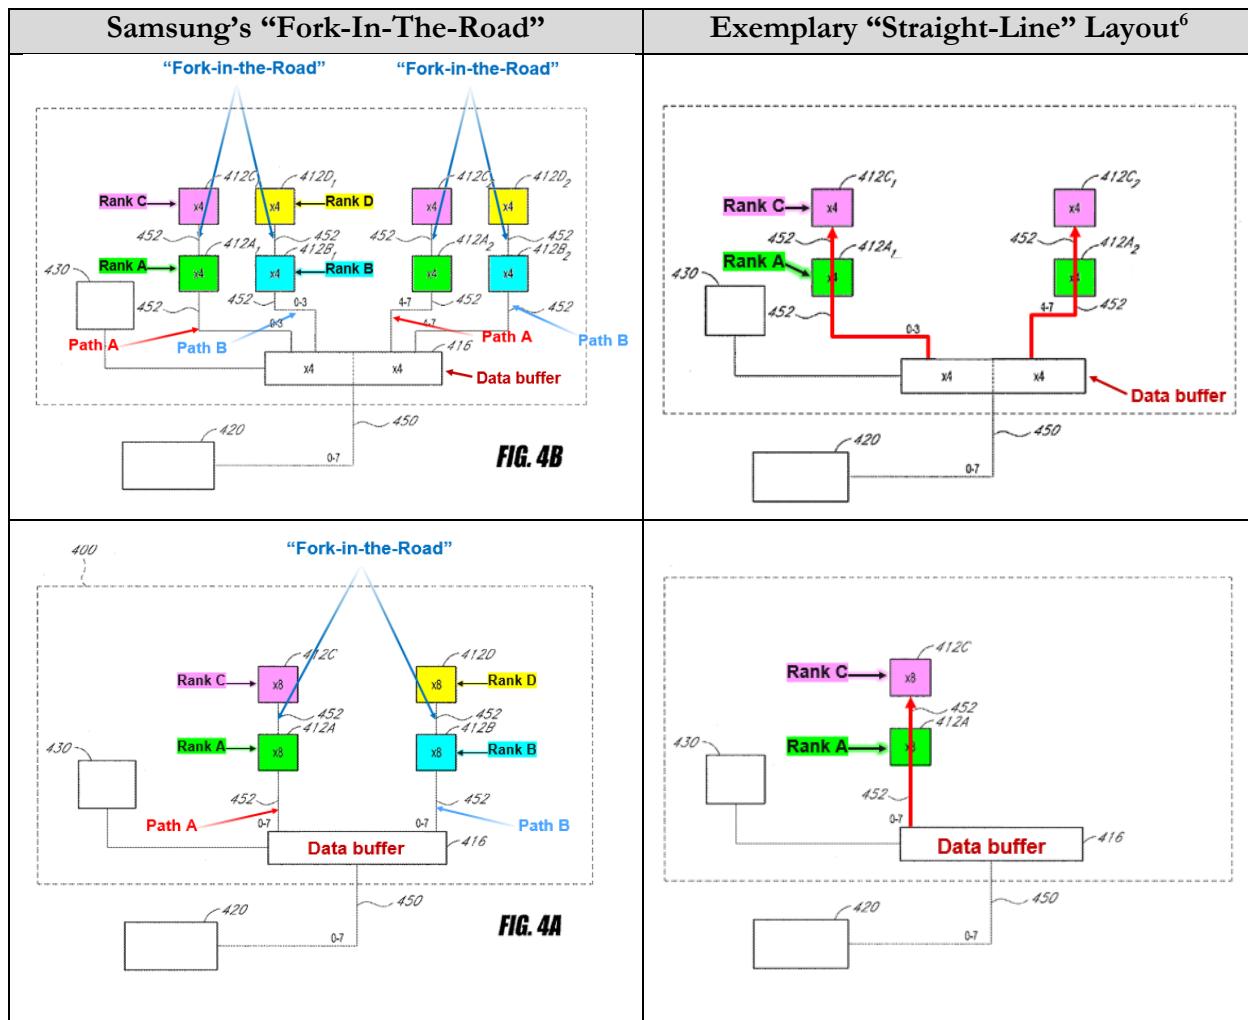
One or more of the data transmission circuits 416, in accordance with an embodiment of this disclosure, is operatively coupled to a corresponding **one or more** of the data lines 452 connected to one or more memory devices 412 in each of the ranks A, B, C, D. For example, in certain embodiments, each data transmission circuit 416 is connected to **one or more** data lines 452 connected to one corresponding memory device in each of the ranks (e.g., memory devices 204A, 204B, 204C, and 204D, as shown in FIG. 3A). Ex. 1 ('339), 14:34-43.

When the data buffer 416 is "operatively coupled to **one** ... of the data lines 452," that is a "Straight Line" configuration. For instance, Figure 4B of the '339 patent shows a "Fork-in-the-Road" configuration where each of the x4 section of the buffer 416 is connected to two data lines 452, with two ranks of memories along each data line 452 (*i.e.*, ranks A & C on the left and ranks B & D on the

---

<sup>5</sup> Samsung argued in its IPR petition against the '339 patent that supposedly "each claim limitation of the 339 Patent is substantially similar to a limitation of the '907 Patent." *E.g.*, Ex. 10 (SS '339 IPR Pet.) at 44. The '907 patent's claims do not require a "Fork-in-the-Road" configuration.

right). But the '339 patent also contemplates that each x4 section is operatively coupled to only one data line 452. *Id.*, 14:34-43 (disclosing data buffer 416 “operatively coupled to a corresponding *one* or more of the data lines 452 connected to one or more memory devices 412 in each of the ranks A, B, C, D”). When a data buffer (or a section of the data buffer) is coupled to only one data line, that would produce a “Straight Line” configuration with *e.g.*, only ranks A and C (or B & D) along the same data path 452 (*see* annotations below).



Furthermore, the load reduction benefit of the claimed distributed buffer arrangement also

<sup>6</sup> Modified Figure 4B showing only two ranks of 'x4' memory devices (A and C) connected to each x4 section of the buffer, while omitting the two others (B and D).

exists in the “Straight Line” configuration. For instance, the ’339 patent teaches that:

The data transmission circuits 416 present a load on the data lines 518 from the write buffer 503 and the read buffer 509. The write buffer 503 is comparable to an input buffer on one of the memory devices 412, and the read buffer 509 is comparable to an output buffer on one of the memory devices 412. Therefore, the ***data transmission circuits 416 present a load to the memory controller 420 that is substantially the same as the load that one of the memory devices 412 would present....*** Ex. 1 (’339), 16:45-54.

In other words, the ’339 patent provides that the load presented by the memory devices 412 to the memory controller 420 can be reduced by virtue of the data buffers 416 in the data path(s), without the selective switching of data paths that is the hallmark of the “Fork-In-The-Road” configuration that Samsung attempts to import. For example, in a “Straight-Line” arrangement of multiple ranks, the memory controller 420 will only see a single load at each data buffer 416. *Id.*, 15:13-18, 16:45-54. For these reasons, the claims are not limited to the “Fork-in-the-Road” configuration.

### B. The Claims Do Not Require Rank Multiplication (’339 claims 1, 11, 19, 27)

Claim 1, and Samsung’s proposed construction, is representative of the parties’ dispute.

Term	Netlist	Samsung
“a module controller ... configurable to receive from the memory controller via the address and control signal lines input address and control signals for a memory write operation to write N-bit-wide write data from the memory controller into a first N-bit-wide rank of the multiple N-bit-wide ranks, and to output registered address and control signals in response to receiving the input address and control signals, wherein the registered address and control signals cause the first N-bit-wide rank to perform the memory write operation by receiving the N-bit-wide write data, wherein the module controller is further configurable to output module control signals in	No construction is necessary (i.e., plain and ordinary meaning).	“a control circuit configurable to receive from the memory controller via the address and control signal lines input address and control signals for a memory write operation to write N-bit-wide write data from the memory controller into a first N-bit-wide rank of the multiple N-bit-wide ranks <i>and corresponding to a number of ranks of memory devices lower than the physical number of ranks of memory devices on the module</i> , and in response to receiving the input address and control signals, to output registered address and control signals <i>corresponding to the number of physical ranks of memory devices on the module</i> , wherein the registered address and control signals cause the first N-bit-wide rank to perform the memory write operation by receiving the N-bit-wide write data, wherein the module controller is further configurable to output module control

response to at least some of the input address and control signals” (’339, claim 1)		signals in response to at least some of the input address and control signals”
---	--	--

Each of the independent claims of the ’339 patent provides that the “module controller” is configurable to receive and output certain signals. Samsung does not assert that the claim language presents any ambiguity, but proposes to add two limitations to each of the “module controller” clauses: (1) that the “input address and control signals” received by the module controller must “correspond[] to a number of ranks of memory devices **lower than** the physical number of ranks of memory devices on the module,” and (2) that the “registered address and control signals” output by the module controller must “correspond[] to the number of physical ranks of memory devices on the module.” In effect, Samsung seeks to import into the claims the concept of “rank multiplication,” a limitation not required by the claim language and not supported by the ’339 specification.

As background, in rank multiplication, an on-module logic simulates a virtual memory module operating as having a second number of (logical or virtual) ranks (e.g. n ranks) of memory devices, when in reality, the memory module has a first number of (physical) ranks of memory devices greater than the first (e.g. 2n). Ex. 13 (’386 patent), 6:63-7:29. This is described in U.S. Pat. No. 7,289,386 (incorporated by reference in the ’339 patent) as one way to expand the number of physical ranks on a memory module. Rank multiplication expands the memory capacity of a system with a limited number of chip-select signal inputs (e.g., 2 which allowed selection among two ranks of memories) by allowing the on-module logic to generate additional chip-select signals to select the additional physical ranks. *Id.*, 2:34-42, 6:63-7:29. That is, rank multiplication helps overcome the limited number of chip selects that a host could issue, but is not needed if the host can already issue an adequate number of chip selects to select among the physical ranks of memory. Ex. 1 (’339), 2:34-42 (“During operation, the ranks of a memory module are selected or activated by control signals that are received from the processor.... Most computer and server systems support one-rank and two-rank memory modules.”).

Nothing in the claim language supports Samsung's blatant claim redrafting. For instance, a comparison of the "module controller" limitation at issue here with other claims that do claim rank multiplication illustrates that the terms at issue here do not require Samsung's imported limitation:

'339 Patent, claim 1	'386 Patent, claim 1
"a module controller ... configurable to receive from the memory controller via the address and control signal lines input address and control signals for a memory write operation to write N-bit-wide write data from the memory controller into <b>a first N-bit-wide rank of the multiple N-bit-wide ranks</b> , and to output registered address and control signals in response to receiving the input address and control signals, wherein the registered address and control signals cause the <b>first N-bit-wide rank</b> to perform the memory write operation by receiving the N-bit-wide write data, wherein the module controller is further configurable to output module control signals in response to at least some of the input address and control signals"	"a logic element ... receiving a set of input control signals from the computer system, the set of input control signals corresponding to a second number of memory devices <b>smaller than the first number of memory devices</b> , the logic element generating a set of output control signals in response to the set of input control signals, the set of output control signals corresponding to the first number of memory devices, wherein the plurality of memory devices are arranged in a <b>first number of ranks, and the set of input control signals corresponds to a second number of ranks of memory modules, the second number of ranks less than the first number of ranks, ....</b> " <sup>7</sup>

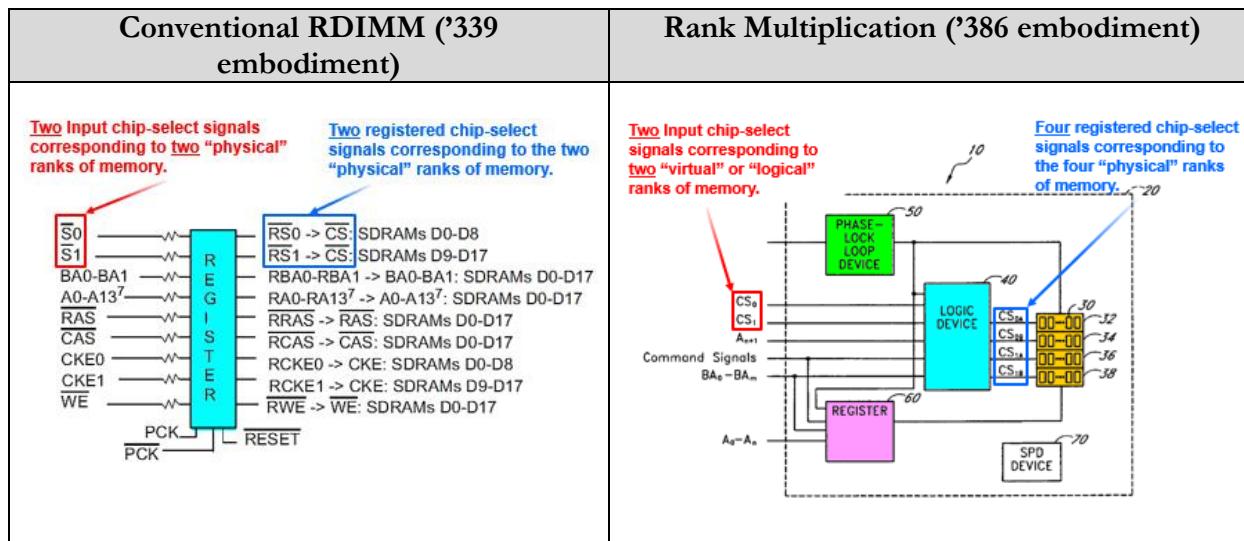
The specification also does not support Samsung's improper attempt to import limitations into the claims. To be sure, the specification, including the disclosures from two incorporated-by-reference patents, describes embodiments compatible with rank multiplication. *See* Ex. 1 ('339), 10:48-50 ("The control circuit 430,430' **may** produce additional chip-select signals or output enable signals based on address decoding.").<sup>8</sup>

But the specification also discloses embodiments where the first and second number of chip-select signals are the same, and hence the number of logical ranks and the number of physical ranks are the same. This obviates the need for rank multiplication. For example, in one such embodiment

<sup>7</sup> The first and second number of memory devices correspond respectively to the number of logical ranks and the number of physical ranks.

<sup>8</sup> The incorporated patents provide "[e]xamples of circuits that can serve as a control circuit 430, 430' [the claimed module controller]" (Ex. 1 ('339), 10:50-53), and do not impose a requirement that the claims of the '339 patent require rank multiplication.

“[t]he control circuit ... registers signals from the control lines 440, 440' in a manner functionally comparable to the address register of a conventional RDIMM.” *Id.*, 10:38-41. As illustrated below, the register of a conventional JEDEC-specified RDIMM receives as input two chip-select signals (S0 and S1) corresponding to two (logical) ranks and outputs two registered chip-select signals (RS0 and RS1) corresponding to two (physical) ranks. *E.g.*, Ex. 14 (JEDEC DDR2 RDIMM) at p. 12-13, 15-16.



As another example, the specification teaches that “embodiments with less than four ranks (e.g., one rank, two ranks, three ranks) … per memory module 402, 402' may be employed.” Ex. 1 ('339), 9:44-49. A two-rank memory module, for example, does not need rank multiplication because, as the incorporated '386 patent teaches, the host could already provide two chip select signals, allowing selection between the two ranks without a need for additional chip select signals or rank multiplication. Ex. 13 ('386), 2:34-42. Samsung’s construction would exclude these preferred embodiments, and is therefore improper because “[a] claim construction that excludes a preferred embodiment is rarely, if ever correct.” *Kaufman v. Microsoft Corp.*, 34 F.4th 1360, 1372 (Fed. Cir. 2022); *see also Oatey Co. v. IPS Corp.*, 514 F.3d 1271, 1276-77 (Fed. Cir. 2008) (“We normally do not interpret claim terms in a way that excludes embodiments disclosed in the specification.”).

**C. The Court Should Resolve the Parties’ Dispute over the Plain Meaning of “time period in accordance with a latency parameter”**

Term	Netlist	Samsung
“time period in accordance with a latency parameter” '339, claims 1, 11, 34, 35	“a time period wherein both the start of the time period and duration of the time period depends on at least a latency parameter”	Plain and ordinary meaning.

This term appears in the clauses defining the time period for which the data paths in the data transmission circuits/buffers are enabled to drive data signals during a command, for example:

each respective byte-wise buffer further includes logic configurable to control the byte-wise data path in response to the module control signals, wherein the byte-wise data path is enabled for a first ***time period in accordance with a latency parameter*** to actively drive a respective byte-wise section of the N-bit wide write data associated with the memory operation from the first side to the second side during the first time period ... Ex. 1 ('339), cl. 1.

*See also, id.*, cls. 11, 34-35 (requiring a first or second set of data buffers to be enabled for a first or second “time period in accordance with a latency parameter”).

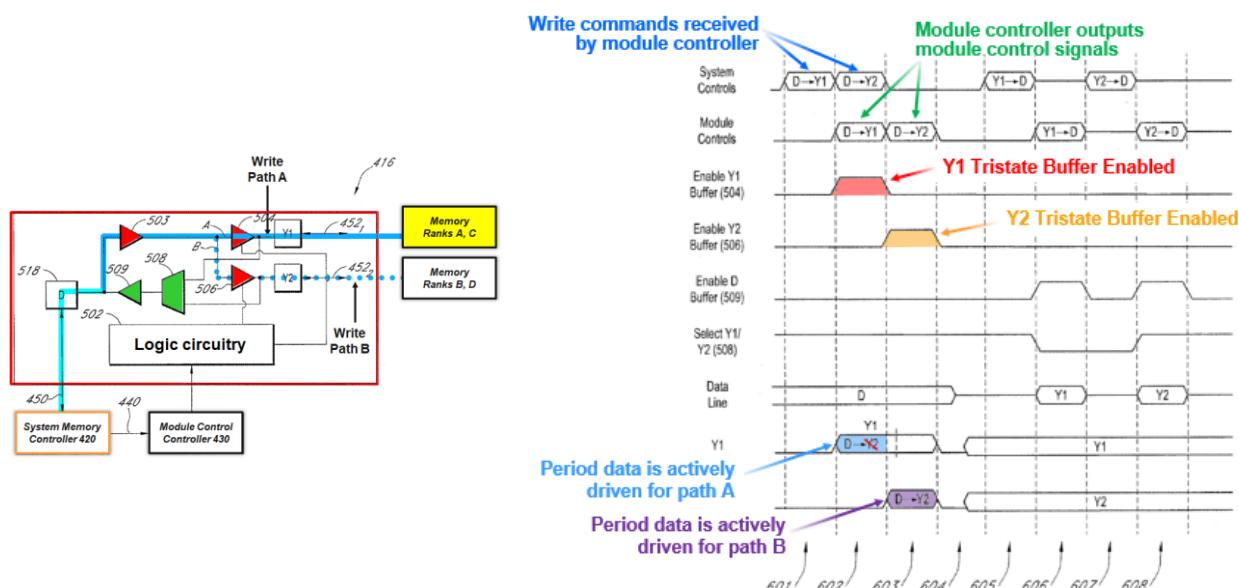
Netlist agrees that the term should be accorded its plain and ordinary meaning. The problem is that Samsung refused to provide its interpretation of the plain and ordinary meaning of the term. Samsung contends in the parallel IPR proceeding, however, that the term requires only the start of the time period to depend on the latency parameter, but not the duration.

The specification teaches that the data transmission to the data buffers may be controlled according to a latency, for example, a Column Address Strobe (CAS) latency. A CAS latency is “a delay time which elapses between the moment the memory controller 420 informs the memory modules 402 to access a particular column in a selected rank or row and the moment the data for or from the particular column is on the output pins of the selected rank or row.” *Id.*, 15:61-16:1. The specification goes on to note that “[d]uring the latency, address and control signals pass from the memory controller 420 to the control circuit 430 which produces controls sent to the control logic circuitry 502 (e.g., via lines 432) which then controls operation of the components of the data

transmission circuits 416.” *Id.*, 16:1-16:6. Thus, “[t]he latency may be used by the memory module to control operation of the data transmission circuit.” *Id.*, 15:65-16:1. As one example, the specification discloses the use of a latency to control the timing of a write operation:

For a write operation, during the CAS latency, the control circuit 430, ..., provides enable control signals to the control logic circuitry 502 of each data transmission circuit 416, whereby the control logic circuitry 502 selects either path A or path B to direct the data.... [W]hen the control logic circuitry 502 receives, for example, an ‘enable A’ signal, a first tristate buffer 504 in path A is enabled and actively drives the data value on its output, while a second tristate buffer 506 in path B is disabled with its output in a high impedance condition. *Id.*, 16:7-25.

In accordance, Fig. 6 shows that in time period 601, the host issues a first write command targeting one of ranks A and C coupled to Y1 terminal (along path A); and in time period 602, the module controller outputs corresponding module control signals to indicate a write operation to a rank coupled to Y1 terminal. In response, the data buffer logic enables the tristate buffer 504 coupled to the Y1 terminal but disables the tristate buffer 506 coupled to the Y2 terminal, so that the data may be written to one of ranks A and C coupled to Y1 terminal, but not to rank B or D coupled to Y2 terminal. *Id.*, 17:66-18:17. At the end of period 602, tristate buffer 504 is disabled. *See, id.*, Fig. 6.



Also in period 602, a second write command is received from the host targeting rank B or D coupled to Y2 terminal. In period 603, the module controller issues corresponding module control signals corresponding to a write operation to a rank coupled to Y2 terminal. In response, the data buffer logic enables the tristate buffer 506 coupled to the Y2 terminal and disables the tristate buffer 504 coupled to the Y1 terminal, so that the data may be written to one of ranks B and D, but not to rank A or C. *Id.*, 18:12-24. The tristate buffer 506 is disabled at the end of period 603. *Id.*, 17:55-57.

Thus, data path A associated with tristate buffer 504 is enabled during period 602, during which the data corresponding to the first write command is passed to Y1 terminal. Similarly, data path B associated with tristate buffer 506 is enabled during period 603, during which the data corresponding to the second write command is passed to Y2 terminal. The start of the enablement time period for data path A (or B) would be affected by, for example, the latency associated with the first (or second) write command because it needs to ensure that data path A (or B) (which includes a write buffer 503 in addition to tristate buffer 504 (or 506)) is enabled in time to allow the first bit of data associated with the first write command to be driven from the host side to the targeted rank coupled to Y1 (or Y2) terminal. The end of the enablement period is also affected by the latency because it needs to ensure that the data path A (or B) is not disabled before all the data associated with the first (or second) write commands propagates through the data buffer and is properly received by the memory devices of the targeted rank. Ex. 1 ('339), 15:61-66 (CAS latency dictates "a delay time which elapses between the moment the memory controller 420 informs the memory modules 402 to access a particular column in a selected rank or row and the moment the data for or from the particular column is on the output pins of the selected rank or row"); *see also, e.g.*, 17:63-65 ("Recalling the CAS latency ..., each write operation extends over two time periods in a pipelined manner."); Ex. 13 ('386), 23:45-55 (discussing the need for pre-amble and post-amble time intervals between data bursts in order to "facilitate good signal integrity"). Both the start and the duration (*i.e.*, the interval between start and

end) is thus in accordance with a latency parameter, consistent with Netlist's proposal.

This interpretation is also consistent with the extrinsic evidence. For example, the Wiley Electrical and Electronics Engineering Dictionary (2004), defines "latency" to include reference to both the "triggering of an action" (i.e., its start) and the "completion" of the action (i.e., its duration). Ex. 15 at 411; *see also* Ex. 16 (Webster's New World, College Dictionary (4<sup>th</sup> Ed., 2000)) at 1071 (period: "the interval between certain happenings [a ten-year period of peace]").

### III. The '506 Patent

The '506 patent also relates to the distributed data buffer architecture implemented in the accused DDR4 LRDIMMs. In conventional memory modules, the "distribution of control signals and a control clock signal in the memory module is subject to strict constraints" Ex. 2 ('506), 2:16-17. The '506 patent describes the need to overcome limitations of prior art memory modules in addressing the complications posed by increasing the number of memory devices on the memory module—and memory density—while maintaining these timing constraints. For example, the '506 patent explains that "read and write leveling" functions in conventional memory controllers "insufficient to insure proper timing of the control and/or data signals received and/or transmitted by the memory modules" as memory operating speed and density continue to increase *Id.*, 2:28-35. To address this need, the '506 patent discloses systems and methods for buffering command, address, and data information via a module control device, a plurality of buffer circuits (also referred to as "isolation devices"), and memory devices on the memory module, resulting in modules with higher capacity and performance.

**A. Claim 14 of the '506 Patent Is Not Limited to “determining a first predetermined amount” During “one or more previous memory operations”**

Term	Netlist	Samsung
“[the method further comprising,] before receiving the input C/A signals corresponding to the memory read operation, determining the first predetermined amount based at least on signals received by the first data buffer”	the step of determining the first predetermined amount based at least on signals received by the first data buffer occurs before the earlier recited step of “receiving, at the module device, input C/A signals corresponding to a memory read operation via the C/A signal lines.”	“during one or more previous memory operations”

Claim 14 recites a method of operating a memory module that includes delaying a read strobe<sup>9</sup> received at a data buffer by a “first predetermined amount” to generate a delayed read strobe for sampling the read data passing through the data buffer to ensure proper timing of data transmission through the module. *See Ex. 2 ('506), cl. 14.* The method recites a step of “receiving, at the module control device, input C/A signals corresponding to a memory read operation via the C/A signal lines” (“Receiving Step”). *Id.*, 21:53-55. The method also recites a determining step of “determining the first predetermined amount based at least on signals received by the first data buffer” “before receiving the input C/A signals corresponding to the memory read operation at the module control device” (“Before Clause”). *Id.*, 22:12-16. The plain language leaves no doubt that the step of “receiving …” in the Before Clause refers to the earlier recited Receiving Step. Thus, all that is required by the disputed term is that the recited “determining” step occurs before the Receiving Step.

Samsung argues, however, that the Before Clause should be interpreted as “during one or more previous memory operations,” in an apparent attempt to redraft claim 14 so that it is identical to claim 1. *Compare, Ex. 2 ('506), 22:11-16 with id., 19:53-55; Dkt. 70 at 2* (parties agreeing to construe

---

<sup>9</sup> A data strobe signal, *e.g.* a read strobe signal, is generally a signal that indicates that another signal, *e.g.*, data or command, is present and valid, and can be used to accurately time the transmission of data throughout the module. Ex. 17 (Jacob) at 318 n.1.

“one or more previous operations” in claim 1 as “one or more previous memory operations”).<sup>10</sup> But the inventors clearly chose to draft claim 14 using the Before Clause, instead of reciting “during one or more previous operations” as in claim 1 and dependent claims 15-16.<sup>11</sup> That choice should be respected. *See Virnetx, Inc. v. Cisco Sys., Inc.*, 767 F.3d 1308, 1316-17 (Fed. Cir. 2014) (rejecting construction of the term “domain name” that read in “hierarchical” limitation because “[t]he specific limitation of hierarchical formatting in the dependent claims strongly suggests that the independent claims contemplate domain names **both with and without** the hierarchical format”). The specification also lends no support to Samsung’s proposed construction, because it does not suggest that the determination must be done during one or more previous memory operations. *See, e.g.*, Ex. 2 (’506), 3:29-34, 4:9-19, 10:11-21, 15:27-16:24 (all examples and not mandates). Rather, these are just preferred “embodiments.” *Id.* Thus, the Court should reject Samsung’s proposed construction.

#### IV. The PMIC Patents (’918/’054)

The ’918 and ’054 patents are directed to a novel power management integrated circuit (“PMIC”) that is mounted on a memory module. DIMMs are typically characterized by, *inter alia*, the generation of the JEDEC standard with which the DIMM complies (e.g., DDR5, DDR4, DDR3). They are mounted on motherboards. Prior to DDR5 DIMMs, a computer system provided power management from the motherboard, external to the memory module. Ex. 18 at 1. The inventions of the PMIC patents effectively enable a memory designer to bring the PMIC and other power management functions onto the module itself, allowing for more precise and accurate regulation of voltages and more efficient power management. As Samsung acknowledges, on-module power-management marks a “major design improvement” that provides for “increased compatibility and

---

<sup>10</sup> Regarding “memory operation” in the agreed-to construction, Netlist’s position is that the operation need not necessarily result in data being actually written to or read from a DRAM.

<sup>11</sup> *Eg.*, Ex. 2 (’506), 22:32-34 (cl. 15: “wherein the second predetermined amount is determined based on signals received by the second data buffer **during one or more previous operations**”).

signal integrity, and providing a more reliable and sustained performance.” *Id.*

#### A. The Multiple “regulated voltages” And “voltage amplitude” Need Not Have Different Voltage Levels

Term	Netlist	Samsung
“first” / “second” / “third” / “fourth” “regulated voltages” ’918, all claims.	Plain and ordinary meaning (that is, first, second, third and fourth voltages that are adjusted, within tolerance, to a particular voltage level).	“first regulated voltage <i>that is distinct from</i> the second, third, and fourth regulated voltages” / “second regulated voltage <i>that is distinct from</i> the first, third, and fourth regulated voltages” / “third regulated voltage <i>that is distinct from</i> the first, second, and fourth regulated voltages” / “fourth regulated voltage <i>that is distinct from</i> the first, second, and third regulated voltages”
“first” / “second” / “third” / “fourth” “voltage amplitude” ’918, all claims.	Plain and ordinary meaning to a (that is, first/ second/ third/ fourth amplitude of voltage which need not all be different).	“first voltage amplitude <i>that is distinct from</i> the second, third, and fourth voltage amplitudes” / “second voltage amplitude <i>that is distinct from</i> the first, third, and fourth voltage amplitudes” / “third voltage amplitude <i>that is distinct from</i> the first, second, and fourth voltage amplitude” / “fourth voltage amplitude <i>that is distinct from</i> the first, second, and third voltage amplitude”
“at least three regulated voltages” ’054: claims 1-15	Plain and ordinary meaning (that is, three or more regulated voltages).	“at least three <i>distinct</i> regulated voltages”
“plurality of regulated voltages” ’054: claims 16, 24	Plain and ordinary meaning (that is, multiple regulated voltages).	“plurality of <i>distinct</i> regulated voltages”

Netlist agrees that there needs to be the recited number of “regulated” voltages present in the accused device or prior art references. The parties dispute whether the claimed “regulated voltages” requires the voltage level (or voltage amplitude) of every “regulated” voltage be different from one another.<sup>12</sup> It does not. During meet-and-confer, Netlist sought clarification on what Samsung meant

---

<sup>12</sup> Netlist does not believe the term “regulated voltage” needs to be construed; but if necessary, Netlist proposes construing it as a “voltage that is maintained, within tolerance, at a particular voltage level.” This is consistent with how a POSITA would understand the term “regulate.” See, e.g., Ex. 15 (Wiley 2004) at 649 (definition of “regulate”: “To maintain a voltage, current, or the like, within specified

by “distinct,” and Samsung drew an analogy to claims directed to a chair with four legs. When asked whether Samsung regarded the chair as having four “distinct” legs if all four legs had the same length, Samsung declined to provide a clear answer.

Samsung’s IPR petitions on the PMIC patents, however, make it clear that Samsung interprets the claims of the PMIC patents as only encompassing implementations in which each recited regulated voltage has a different voltage amplitude (*i.e.*, voltage level). *See, e.g.*, Ex. 20 (IPR2022-00999, Paper 1) at 27 (concerning ’054 claim 1 which does not recite “voltage amplitude,” Samsung arguing “Grounds 1A-1B have two **voltage amplitudes** that are the same, … while Ground 1C does not, consistent with a *narrower* interpretation”). Samsung is wrong: (1) with respect to claims that only recite “regulated voltages,” Samsung’s “narrower” interpretation injects an unrecited term “voltage amplitude” into the claims; and (2) even for claims that recite “voltage amplitude,” Samsung’s construction amounts to importing limitations from a preferred embodiment. Both are improper.

First, the claims differentiate between “voltage” on the one hand and “voltage amplitude” on the other. For example, certain claims expressly both voltage and voltage amplitudes. For instance, claim 1 of the ’918 patent recites “*a first/second/third/fourth regulated voltage having a first/second/third/fourth voltage amplitude.*” Ex. 3 (’918), 38:25-32; *see also id.*, 39:65-67 (“wherein the first, second, third and fourth regulated voltages have first, second, third, and fourth voltage amplitude, respectively”). Other claims, however, recite only “voltage[s],” and not “voltage amplitude[s].” *See, e.g.*, Ex. 3 (’918), 41:4-8 (“… provide the first, second and third regulated voltages, respectively; and … provide the fourth regulated voltage”); Ex. 4 (?054), 38:26-31 (each buck converter “is configured to produce a regulated voltage of the at least three regulated voltages”), 39:66-40:4 &

---

values”); Ex. 19 (Electronics 3rd Ed.) at 350 (definition of “regulated voltage”: “a voltage that remains steady, unaffected by changes in LOAD current or mains supply voltage”); Ex. 16 (Webster’s 4th) at 1207 (“**2** [t]o adjust to a particular standard, rate, degree, amount, etc. [*regulate* the heat] **3** to adjust (a clock, etc.) so as to make operate accurately”).

41: 9-13 (each buck converter “is configured to produce a regulated voltage of the plurality of regulated voltages”). These claims do not recite “voltage amplitudes” and no such limitation should be imposed on this absent term. That is, the inventors clearly knew how to draft the claims to include limitations on voltage amplitudes when desired, and it is improper for Samsung to read voltage amplitude limitation into the claims when the inventors chose to omit that language from the claims. *Phillips v. AWH Corp.*, 415 F.3d 1303, 1314 (Fed. Cir. 2005) (“Differences among claims can also be a useful guide in understanding the meaning of particular claim terms.”).

Second, the claims as a whole make clear that the voltage amplitudes need not all be different. For example, the ’918 patent, claim 1 requires that “a first one of the second and fourth voltage amplitudes is less than a second one of the second and fourth voltage amplitudes.” Ex. 3 (’918), 38:49-52. This limitation requires that one of the second and fourth voltage amplitudes be less than the other one, *i.e.*, the second and fourth voltage amplitudes differ from each other. If each voltage amplitude had to be different from the others by virtue of being recited as “first,” “second,” “third” and “fourth,” then this express limitation on the relative amplitude would be superfluous. Likewise, dependent claim 2 of the ’054 patent recites: “wherein the first regulated voltage has a first voltage amplitude, and the second regulated voltage has a second voltage amplitude; and wherein a first one of the first and second voltage amplitudes is less than a second one of the first and second voltage amplitudes.” Ex. 4 (’054), 38:45-50. Again, if first and second voltage amplitudes had to be different, there would be no point to further specify that one of the first and second voltage amplitudes is less than the other. *Apple, Inc. v. Ameranth, Inc.*, 842 F.3d 1229, 1234 (Fed. Cir. 2016) (“Construing a claim term to include features of that term already recited in the claims would make those expressly recited features redundant.”). As another example, dependent claim 29 of the ’918 patent recites “wherein the plurality of SDRAM devices are configured to receive **at least one of** the first, second, third and fourth regulated voltages having a voltage amplitude of 1.8 volts.” Ex. 3 (’918), cl. 29. By reciting “at least one,” this claim

encompasses receiving two, three or four voltages of 1.8V—that is, voltages at the same level—from the recited set of regulated voltages. This contradicts Samsung’s proposed construction.

Third, nothing in the specification narrows the claim scope in the manner proposed by Samsung. For example, in one preferred embodiment, volatile memory 1032 and nonvolatile memory 1042 are both powered by a 1.8V first voltage 1102 produced by the buck-converter of sub-block 1122. Ex. 3 (’918), 29:39-44, Fig. 16. But in another embodiment, “the volatile memory elements 1032 and nonvolatile memory elements 1042 are powered using independent voltages and are not both powered using the first voltage 1102.” *Id.*, 29:61-64. But if volatile memory 1032 and nonvolatile memory 1042 are both supposed to operate at 1.8V, then even if they are powered by independent voltages, those independent voltages need to each provide the 1.8V voltage level. This suggests that independent voltages need not have different voltage levels.

Samsung’s position is unsupported even assuming *arguendo* that the embodiment described in 29:39-54 of the ’918 patent was the only embodiment describing multiple voltages generated by conversion element 1120. The Federal Circuit has held “[e]ven when the specification describes only a single embodiment, the claims of the patent will not be read restrictively unless the patentee has demonstrated a clear intention to limit the claim scope using ‘words or expressions of manifest exclusion or restriction.’” *Hill-Rom*, 755 F.3d at 1371-72 (citations omitted). The record contains no such “manifest exclusion or restriction” of the claimed voltages to ones of all different voltage levels. *See generally* Ex. 3 (’918), 29:55-58 (“Although described with respect to certain example embodiments, one of ordinary skill will recognize from the disclosure herein that the conversion element 1120 may include alternative embodiments.”). Moreover, the specification describes a 1.8V “third voltage” 1108 (an input voltage), which happens to have the same voltage level as the “first voltage” 1102 (an output voltage). *Id.*, 28:15-19, 29:39-44. This again evidences that the inventors did not intend a voltage (*e.g.*, a “first” voltage) to necessarily have a different voltage level than a differently numbered voltage (*e.g.*,

a “third” voltage). *Id.*

Samsung’s citation to IPR proceedings and prosecution history of other Netlist patents does not demonstrate otherwise. *See* Dkt. 70-2 at 7 (citing, *inter alia*, numerous Netlist submissions from IPRs for the ’833 and ’831 patents). First, those claims did not even recite “voltage” or “voltage amplitude.” *See, e.g.*, Ex. 21 ’833, cl. 1 (reciting first/second/third clock frequencies); Ex. 22 (14/840,865) at 11 (first/second volatile memory subsystems); Ex. 25 (15/934,416) at 9 (“first and second commands”). Second, even there, the applicants did not argue that “first” and “second” cannot be of the same type or value unless the claims required otherwise. For instance, ’833 claim 1 recites the “third clock frequency being less than the first clock frequency.” Ex. 21, cl. 1. As such, when the applicants argued that a prior art reference did not disclose different operation frequencies, they were only arguing that the references did not disclose the required elements. *E.g.*, Ex. 23 (10/18/2011 OA) at 14-15; Ex. 24 (5/24/2012 OA) at 7-8; *see also, e.g.*, Ex. 22 (11/6/2017 OA) at 11 (applicants argued that the first and second volatile memories should not be the same DRAM, not that they cannot both be DRAMs); Ex. 25 (3/12/2020 OA) at 9-11 (applicants did not preclude the possibility that the first and second commands can both be, *e.g.*, a read or write command; and only pointed out that the prior art did not distinguish between whatever the first and second commands the office mapped to and failed to disclose the limitation).

#### B. Samsung’s Proposed Construction For “A Second Plurality of Address And Control Signals” Confuses Rather Than Clarifies the Meaning of the Term

Term	Netlist	Samsung
“a second plurality of address and control signals”	Plain and ordinary meaning (that is, a second set of address and control signals).	“a second plurality of address and control signals <i>that are distinct from</i> a first plurality of address and control signals”

During meet-and-confer, Samsung indicated that it would be open to construing the term as “a second set of address and control signals distinct from the first.” As noted above, Samsung would not explain what it meant by “distinct.” Given Samsung’s refusal to articulate what it means by

“distinct,” Samsung’s proposed construction just creates more confusion. This is anathema to the claim construction process, and Samsung’s proposed construction should therefore be rejected.

There is also no need to construe the term. The claim itself provides all that is needed: “the at least one circuit operable to (i) receive a first plurality of address and control signals via the first portion of the plurality of edge connections, and (ii) output a second plurality of address and control signals to the plurality of SDRAM devices.” Ex. 3 (’918) at 38:43-47. That is, the claim makes clear that the recited “second plurality of address and control signals” are the signals output from the at least one circuit, as opposed to the first “plurality” of signals that are received by the circuit. The claim, consistent with the specification, imposes no further restrictions, and does not preclude the input and output signals be wholly or partially the same. *E.g.*, Figs. 12-15; *see also supra*, Part IV.A. (explaining why Samsung’s extrinsic evidence on completely different claim elements is irrelevant).

### C. A “dual buck converter” Need Not Produce Two “Distinct” Voltages

Term	Netlist	Samsung
“dual buck converter” / “dual-buck converter”	Plain and ordinary meaning (that is, a buck converter with two regulated voltage outputs whose amplitude may be the same or different).	“buck converter with two outputs outputting two <i>distinct</i> regulated voltages”

These terms are in certain dependent claims of the PMIC patents, which recite configurations in which two of the claimed buck converters are “configured to operate as a dual buck converter.” Ex. 3 (’918), cl. 2, 17, 28; Ex. 4 (’054), cl. 15. Netlist agrees with Samsung that a “dual buck converter” provides two regulated voltage outputs, but disagrees that those outputs need to be “distinct” or have different voltage levels.

The specification discloses a “dual buck converter” 1124 (Fig. 16) that provides two regulated voltages 1104 and 1105. Ex. 3 (’918), Fig. 16, 29:18-31 (depicting sub-block 1124 as comprising a dual buck converter outputting two voltages 1104 and 1105). “In one *example embodiment*,” the voltages 1104 and 1105 have different voltage levels (2.5V and 1.2V respectively). *See, e.g., id.*, 29:46-

49. But this is just “one example embodiment.” Indeed, the specification emphasizes that “alternative embodiments” “may produce different voltage values” than that disclosed in the above example. *Id.*, 29:55-61. Nothing in the PMIC patents mandates that the two “regulated voltage outputs” by the dual-buck converter be different. *Hill-Rom*, 755 F.3d at 1371-72 (Federal Circuit “has expressly rejected the contention that if a patent describes only a single embodiment, the claims of the patent must be construed as being limited to that embodiment”).

#### D. The Preamble of the PMIC Patent Claims Is Limiting

Term	Netlist	Samsung
“memory module” (all claims)	The preamble is limiting.	The preamble is non-limiting.

The preamble of the PMIC patent claims is limiting because it provides antecedent basis for a term in the body of the claims. The preamble of the claims of the ’918 patent and claims 1-15 and 25 recites “[a] memory module.”<sup>13</sup> This provides the antecedent basis for the later recited “memory module” in the body of the claims. *E.g.*, Ex. 3 (’918) at 38:18-24 (PCB interface “including a plurality of edge connections configured to couple power, data, address and control signals between the **memory module** and the host system”), 39:53-59 (same), 40:50-56; Ex. 4 (’054) at 38:19-25, 42:3-6. Further, the parties agree that the terms “operable state,” “first operable state,” and “second operable state” in claims 4-7, 11-12, 16-17, 23, and 25-26 of the ’054 patent refer to a “state in which the **memory module** is operated,” *see* Dkt. 70 at 2, confirming that the term “memory module” is necessary to understand the scope of the claim. *Catalina Mktg. Int’l, Inc. v. Coolsavings.com, Inc.*, 289 F.3d 801, 808 (Fed. Cir. 2002) (“[A] preamble limits the invention if it recites essential structure or steps, or if it is ‘necessary to give life, meaning, and vitality’ to the claim.”).

---

<sup>13</sup> “Memory module” is a term of art and the term is used in common designs such as DIMMs and single in-line memory modules (“SIMM”). *See also* Ex. 3 (’918), 21:24-55.

E. “pre-regulated input voltage” / “input voltage”

Term	Netlist	Samsung
“pre-regulated input voltage” (’918 claims 16, 22, 30)	Plain and ordinary meaning (that is, modulated input voltage).	“regulated voltage generated on the memory module from an input voltage”

The parties agree that a “pre-regulated” input voltage should itself be a regulated voltage.<sup>14</sup>

The parties disagree, however, where that voltage regulation needs to occur. Samsung argues that it must be “generated on the memory module from an input voltage.” Netlist disagrees: the claims only requires that the voltage is a regulated one when received by the respective buck converters or converter circuit. *See Ex. 3 (’918), 39:60-61 (“first, second, and third buck converters configured to receive a pre-regulated input voltage”), 39:63-64 (“a converter circuit configured to reduce the pre-regulated input voltage to provide a fourth regulated voltage”).*

Samsung apparently believes that because the claims recite, for example, converters “configured to receive **a pre-regulated input voltage**” and a voltage monitor circuit that is “configured to monitor **an input voltage** received via a first portion of [a] plurality of edge connections,” (e.g., Ex. 3 (’918), 39:60-62, 40:7-9), the pre-regulated input voltage must be generated from the later appearing “input voltage” that is being monitored. But the claims themselves do not relate the recited “pre-regulated input voltage” and the separate “input voltage.” And while a preferred embodiment in Fig. 16 teaches that pre-regulated voltages can be generated on-board from an input system voltage, that is not required. For instance, Figure 16 “comprises one or more of the components described … with respect to FIG. 12,” including “the second power supply 1080 and the switch 1090.” Ex. 3 (’918), 27:59-28:2. The second power supply 1080 is described as including a capacitor bank with step-up and step-down transformers for serving as a power supply when the

---

<sup>14</sup> The term “input voltage” in “pre-regulated input voltage,” if construed, should receive its plain and ordinary meaning: voltage that is provided to the earlier mentioned converters or converter circuit.

memory module encounters a trigger event (such as power failure). *Id.*, 25:54-62, 26:8-13, Fig. 12. This corresponds to the second power subsystem 1140 of Fig. 16, which also comprises a capacitor array and step-up/step-down regulator that generates a pre-regulated voltage 1112 and is used when a trigger event has occurred. *Id.*, 28:39-43, 28:54-67, 30:35-45.<sup>15</sup> The specification teaches that “the second power supply 1080 may not be on the same printed circuit board 1020 and may be tethered to the printed circuit board 1020[.]” *Id.*, 26:32-35. As such, the specification makes clear that the power subsystem used to generate a pre-regulated voltage need not be “on the memory module from an input voltage [received from the module interface]” as Samsung contends.

## V. The HBM Patents ('060/'160)

The HBM patents disclose systems and methods for optimizing a load in a memory package featuring a control die, array dies, and numerous die interconnects. In contrast to traditional DDR modules in which different DRAM devices are packaged individually and then assembled on a common printed circuit board, the inventions of the HBM patents are directed to a memory package having multiple vertically stacked array dies in electrical communication with a control die. The control die includes conduits with corresponding drivers for driving data, command and/or address signals to the array dies along the die interconnects. The drivers are subject to a load by virtue of being in electrical communication with the corresponding die interconnects and array dies. In conventional memory packages, the die interconnects are in communication with *each* of the array dies, which disadvantageously increases the load on the drivers, requiring larger driver sizes which increase the size of the control die and consumes more power. Ex. 5 ('060), 2:8-15. To address this problem, the HBM patents disclose memory packages with multiple die interconnects in electrical communication

---

<sup>15</sup> In connection with Figures 12-14, the specification describes the condition associated with detecting a trigger condition is likely to occur as the third state and the condition associated with occurrence of a trigger condition as the second state. In connection with Figures 16-17, however, the specification describes the former as the second state and the latter as the third state. Ex. 3 at 25:58-26:3, 28:39-43.

with some, but not all, of the array dies. The invention enables the memory packages to be designed with a smaller footprint, and lowers power consumption by reducing the load on each conduit and the corresponding driver. *Id.* 7:22-8:62, 18:4-13.

#### A. “array die” Does Not Exclude a DRAM Circuit

Term	Netlist	Samsung
“array die” (all claims)	Plain and ordinary meaning (that is, a die including memory cells)	“array die that is different from a DRAM circuit”

The term “array die” should be accorded its plain and ordinary meaning, which is a die including memory cells. The specification supports this construction. For example, the specification teaches that an “array die” needs to include “memory cells.” *See, e.g.*, Ex. 5 ('060), 1:57-62 (“In some cases, the control die 130 **may include memory cells and therefore, also serve as an array die....** Alternatively, the control die 130 and the array dies 110 may be distinct elements and the control die 130 **may not include any memory cells.**”). That a control die having memory cells can also serve as an array die indicates to a POSITA that a fundamental characteristics of an array die is the presence of memory cells on the die. *See also, e.g.*, *id.*, 1:23-29, 1:33-42, 1:63-2:2, 5:17; 1:53-56 (“The array dies 110 are configured to transfer data (e.g. read or write) to or from the selected **memory cells** identified by the command, address, and chip select signals via the die interconnects.”).<sup>16</sup>

In contrast, Samsung has proposed that “array die” should be construed as “array die that is different from a DRAM circuit.” Samsung apparently bases its construction on statements that Netlist made during prosecution of the '060 patent. In particular, during prosecution, to distinguish pending claim 1 (issued as claim 1) from a prior art reference called Rajan, the applicant remarked:

---

<sup>16</sup> To the extent there is any dispute as to the meaning of the term “die,” a POSITA would understand that a “die” refers to “[a] single piece of silicon that contains one or more circuits and is or will be packaged as a unit.” *E.g.*, Ex. 28 (IEEE 100 The Authoritative Dictionary of IEEE Standards Terms (7th Ed. 2000)), at 301.

Rajan merely stacks DRAM circuits 206A-D, which are different from array dies. **As a result**, Rajan's buffer chip 202 also operates very differently from the control die in claim 1. Ex. 26 (1/13/2014 Remarks) at 10.

Samsung is wrong. First, reading the file history as a whole, a POSITA would not regard the above statement as disclaiming DRAM circuits from the scope of claimed array dies. This is because when read in the context of prosecution, the statements were to explain why Rajan's buffer circuit had a different structure, function and operation than the claimed "control die." *See id.* at 10-11. Indeed, there were remarks in the same Office Action response that indicate the inventor did not disclaim "DRAM circuits." Specifically, the applicants remarked that "[t]he arguments regarding claim 1 is applicable to claim 29." Ex. 26 at 14. But claim 29 recites:

a plurality of DRAM packages, each DRAM package comprising:  
a plurality of data terminals via which the DRAM package communicate data with the memory control hub, and a plurality of control terminals to receive the control signals;  
a plurality of array dies arranged in a stack, including a first group of array dies and a second group of at least one array die;

*Id.* at 6; *also compare id.* at 2, cl. 1 *with id.* at 6, cl. 29 (the requirements for each DRAM package in cl. 29 generally track those for the memory package of claim 1). It is well established that a term in the same patent should generally be accorded the same meaning. *In re Varma*, 816 F.3d 1352, 1363 (Fed. Cir. 2016) ("[T]he principle that the same phrase in different claims of the same patent should have the same meaning is a strong one, overcome only if 'it is clear' that the same phrase has different meanings in different claims."). Thus the term "array dies" in claim 29 should be accorded the same meaning as in claim 1. But in claim 29, the array dies are a part of a DRAM package. Ex. 26 at 6. It therefore cannot be the case that array dies in a DRAM package would exclude DRAM circuits. Thus, by maintaining claim 29 and stating that the arguments for claim 1 also applied to claim 29, the applicant made clear that he did not intend to exclude "DRAM circuits" from the scope of "array dies."

Second, disclaimer, if any, would be limited to **Rajan's** DRAM circuits 206A-D **as stacked by Rajan.** *See Genuine Enabling Tech. LLC v. Nintendo Co.*, 29 F.4th 1365, 1374-75 (Fed. Cir. 2022)

(disclaimer of claim scope for “input signal” limited to “signals below the audio frequency spectrum” where applicant repeatedly distinguished his invention on that ground, but “t[o] the extent [applicant’s] statements may implicate other claim scope—such as signals of frequency up to 500 Hz—the record does not rise to the level of establishing a ‘clear and unmistakable’ disavowal”); *Tech. Props. Ltd. LLC v. Huawei Techs. Co.*, 849 F.3d 1349, 1358 (Fed. Cir. 2017)) (no disclaimer where statements are ambiguous or amenable to multiple reasonable interpretations). Rajan’s “DRAM circuits 206A-D” were then available commodity DRAMs. *See* Ex. 27 (Rajan) [0018] (“[I]n various embodiments, one or more of the memory circuits … may include a monolithic memory circuit. For instance, such monolithic memory circuit may take the form of [DRAM]. Such DRAM may take any form including, but not limited to synchronous (SDRAM), double data rate synchronous (DDR DRAM, DDR2 DRAM, DDR3 DRAM, etc.) … and/or any other type of DRAM.”); *id.*, [0043]-[0045].

Third, Samsung has therefore taken the remark on *Rajan* out of context. In the Office Action response, the inventor explained why *Rajan*’s DRAM stack cited by the Office resulted in structural and functional differences between *Rajan*’s buffer die and the claimed control die. *Id.* at 10-11. For example, Netlist explained that according to the cited portions of Rajan, “all that is required of the buffer chip 202 is the capability of buffering the stack of DRAM circuits 206A-D to electrically and/or logically resemble at least one larger capacity DRAM circuit to the host system.” Ex. 26 at 11. As a result, argued the inventor, Rajan did not disclose the claimed “first data conduit,” “second data conduit” or control logics required for the claimed “control die.” *Id.* That is, the remark on *Rajan*’s DRAM circuits is to differentiate the claim requirement from the prior art.

Hence, in summary, the applicant made no “clear and unmistakable disavowal” of DRAM circuits from the scope of “array dies” in his claims. To the contrary, by expressly reciting “array dies” as part of a DRAM package (*see* claim 29), the inventor made clear that he did not intend to categorically disclaim DRAM circuits. Ex. 26 at 6-7. Were that not the case, it would have made no

sense for the inventor to keep a claim directed specifically to DRAM packages with array dies.

**B. A “chip select signal” Is Not Limited to Selecting Only One Array Die**

Term	Netlist	Samsung
“chip select signal”	“signal for enabling or selecting one or more array dies for data transfer”	Plain and ordinary meaning.
“chip select conduits”	“conduits for transmitting” “chip select signals,” as construed above	Plain and ordinary meaning.

These terms appear in claims 6, 11-14, 16-19, 20, 21, 23-28 of the ’060 patent. For example, claim 11 of the ’060 patent recites:

a control die comprising at least a first data conduit between the first die interconnect and a first terminal of the plurality of input/output terminals, at least a second data conduit between the second die interconnect and the first terminal, and ***chip select conduits for providing chip select signals*** to respective array dies; wherein the control die further comprises a control circuit to control respective states of the first data conduit and the second data conduit to drive a data signal to ***an array die selected by at least one of the chip-select signals***. Ex. 5 (’060), 25:16-25.

Netlist’s proposed construction for “chip select signal” follows the specification, which provides that “[e]ach array die 110 also includes a chip select port 144, with the chip select ports 144 of the array dies 110 configured to receive corresponding ***chip select signals to enable or select the array dies for data transfer.***” Ex. 5 (’060), 1:49-56. The claims also supports Netlist’s construction of “chip select conduits” as conduits for transmitting chip select signals to corresponding array dies, as does the specification. *See id.*, cl. 6 (“control die further compris[ing] chip-select conduits, [and] ... third die interconnects coupled between respective chip-select conduits and respective ones of the plurality of stacked array dies.”), 25:22-25 (cl. 11), cl. 13 (“wherein the chip select conduits include drivers to drive the chip select signals to the respective array dies”).

September 2, 2022

Respectfully submitted,

/s/ Jason Sheasby

Samuel F. Baxter  
Texas State Bar No. 01938000  
sbaxter@mckoolsmith.com  
Jennifer L. Truelove  
Texas State Bar No. 24012906  
jtruelove@mckoolsmith.com  
**MCKOOL SMITH, P.C.**  
104 East Houston Street Suite 300  
Marshall, TX 75670  
Telephone: (903) 923-9000  
Facsimile: (903) 923-9099

Jason G. Sheasby (*pro hac vice*)  
jsheasby@irell.com  
Annita Zhong, PhD (*pro hac vice*)  
hzhong@irell.com  
Thomas C. Werner (*pro hac vice*)  
twerner@irell.com  
Yanan Zhao (*pro hac vice*)  
yzhao@irell.com  
Michael W. Tezyan (*pro hac vice*)  
mtezyan@irell.com

**IRELL & MANELLA LLP**  
1800 Avenue of the Stars, Suite 900  
Los Angeles, CA 90067  
Tel. (310) 277-1010  
Fax (310) 203-7199  
*Attorneys for Plaintiff Netlist, Inc.*

**CERTIFICATE OF SERVICE**

I hereby certify that, on September 2, 2022, a copy of the foregoing was served to all counsel of record.

/s/ Jason Sheasby  
Jason Sheasby